



US009311929B2

(12) **United States Patent**
Kroeker et al.

(10) **Patent No.:** **US 9,311,929 B2**
(45) **Date of Patent:** ***Apr. 12, 2016**

(54) **DIGITAL PROCESSOR BASED COMPLEX ACOUSTIC RESONANCE DIGITAL SPEECH ANALYSIS SYSTEM**

7,085,721 B1 * 8/2006 Kawahara et al. 704/258
7,457,756 B1 11/2008 Nelson et al.
7,492,814 B1 2/2009 Nelson

(Continued)

(71) Applicant: **Eliza Corporation**, Danvers, MA (US)

FOREIGN PATENT DOCUMENTS

(72) Inventors: **John P. Kroeker**, Hamilton, MA (US);
Janet Slifka, Hopkinton, MA (US);
Richard S. McGowan, Lexington, MA (US)

KR 1020040001131 A 1/2004
KR 1020050072976 7/2005

(Continued)

(73) Assignee: **Eliza Corporation**, Danvers, MA (US)

OTHER PUBLICATIONS

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 255 days.

Kaniewska, Magdalena, "On the instantaneous complex frequency for pitch and formant tracking", Signal Processing Algorithms, Architectures, Arrangements, and Applications (SPA), Sep. 25-27, 2008, pp. 61 to 66.*

(Continued)

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **13/665,486**

Primary Examiner — Jesse Pullias

(22) Filed: **Oct. 31, 2012**

(74) Attorney, Agent, or Firm — Russ Weinzimmer & Associates, P.C.

(65) **Prior Publication Data**

US 2014/0122067 A1 May 1, 2014

(57) **ABSTRACT**

(51) **Int. Cl.**
G10L 15/02 (2006.01)
G10L 19/02 (2013.01)
G10L 25/15 (2013.01)

A speech analysis system uses one or more digital processors to reconstruct a speech signal by accurately extracting speech formants from a digitized version of the speech signal. The system extracts the formants by determining an estimated instantaneous frequency and an estimated instantaneous bandwidth of speech resonances of the digital version of the speech signal in real time. The system digitally filters the digital speech signal using a plurality of complex digital filters in parallel having overlapping bandwidths to ensure that substantially all of the bandwidth of the speech signal is covered. This virtual chain of overlapping complex digital filters produces a corresponding plurality of complex filtered signals. A first estimated frequency and a first estimated bandwidth is generated for each of the filtered signals, and speech resonances of the input speech signal are identified therefrom.

(52) **U.S. Cl.**
CPC **G10L 25/15** (2013.01)

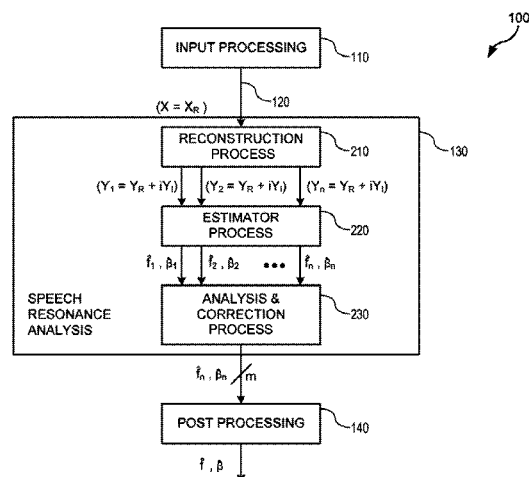
(58) **Field of Classification Search**
USPC 704/200–230
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,346,262 A * 8/1982 Willems et al. 704/217
5,463,716 A * 10/1995 Taguchi 704/209
6,577,968 B2 6/2003 Nelson

38 Claims, 11 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

7,522,594	B2	4/2009	Piche et al.	
7,624,195	B1	11/2009	Biswas et al.	
7,756,703	B2 *	7/2010	Lee et al.	704/209
2004/0228469	A1	11/2004	Andrews et al.	
2005/0049866	A1 *	3/2005	Deng et al.	704/240
2007/0071027	A1	3/2007	Ogawa	
2007/0112954	A1	5/2007	Ramani et al.	
2007/0276656	A1 *	11/2007	Solbach et al.	704/200.1
2008/0082322	A1 *	4/2008	Joublin et al.	704/209

FOREIGN PATENT DOCUMENTS

KR	1020060013152	A	2/2006
KR	10-0731330		6/2007

OTHER PUBLICATIONS

Potamianos et al., "Speech Formant Frequency and Bandwidth Tracking Using Multiband Energy Demodulation", 1995 International Conference on Acoustics, Speech, and Signal Processing, ICASSP-1995, May 9-12, 1995, vol. 1, pp. 784 to 787.*

Jones et al., "Instantaneous Frequency, Instantaneous Bandwidth and the Analysis of Multicomponent Signals", 1990 International Con-

ference on Acoustics, Speech, and Signal Processing, ICASSP-1990, Apr. 3-6, 1990, vol. 5, pp. 2467 to 2470.*

Kenneth N. Stevens, Acoustic Phonetics, Book, 1998, pp. 258-259, Massachusetts Institute of Technology, United States.

Francois Auger and Patrick Flandrin, Improving the Readability of Time-Frequency and Time-Scale Representations by the Reassignment Method, publication, 1995, pp. 1068-1089, vol. 43, IEEE.

T.J. Gardner and M.D. Magnasco, Instantaneous Frequency Decomposition: An Application to Spectrally Sparse Sounds with Fast Frequency Modulations, publication, 2005, pp. 2896-2903, vol. 117, No. 5, Acoustical Society of America, U.S.

Randy S. Roberts, et al., Computationally Efficient Algorithms for Cyclic Spectra Analysis, magazine, 1991, pp. 38-49, IEEE, US.


David T. Blackstock, Fundamentals of Physical Acoustics, book, 2000, pp. 42-44, John Wiley & Sons, Inc., US & Canada.


Iwao Sekita et al., Complex Autoregressive Model and its Properties, publication, 1999, pp. 1-6, Electrotechnical Laboratory, Japan.


Saeed V. Vaseghi, Advanced Digital Signal Processing and Noise Reduction, book, 2006, pp. 213-214, 3rd edition, John Wiley & Sons, Ltd., England.


Malcolm Slaney, An Efficient Implementation of the Patterson-Holdsworth Auditory Filter Bank, technical report, 1993, pp. 2-41, Apple Computer Technical Report #35, Apple Computer Inc., US.


* cited by examiner











310

FIG. 4

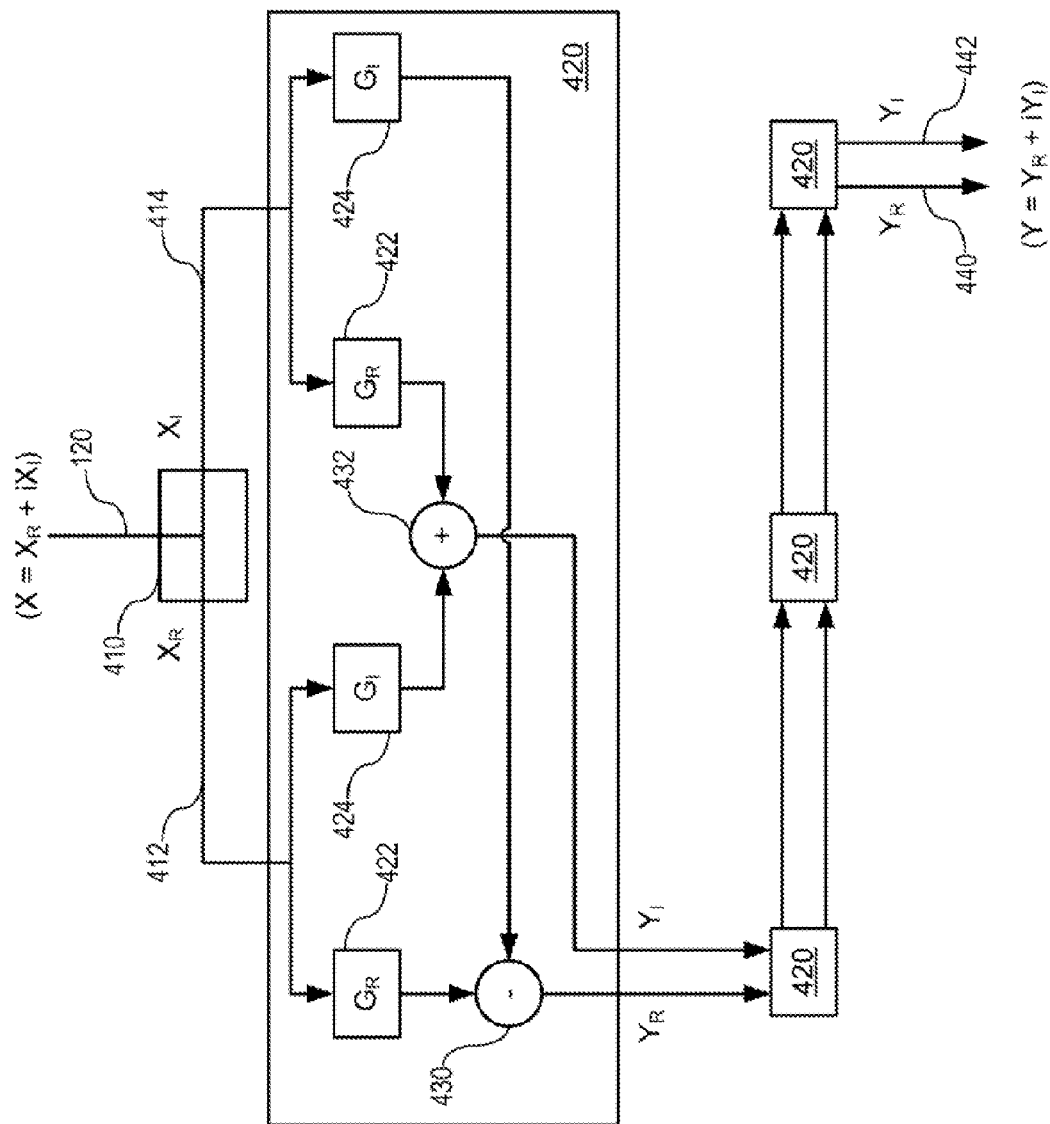


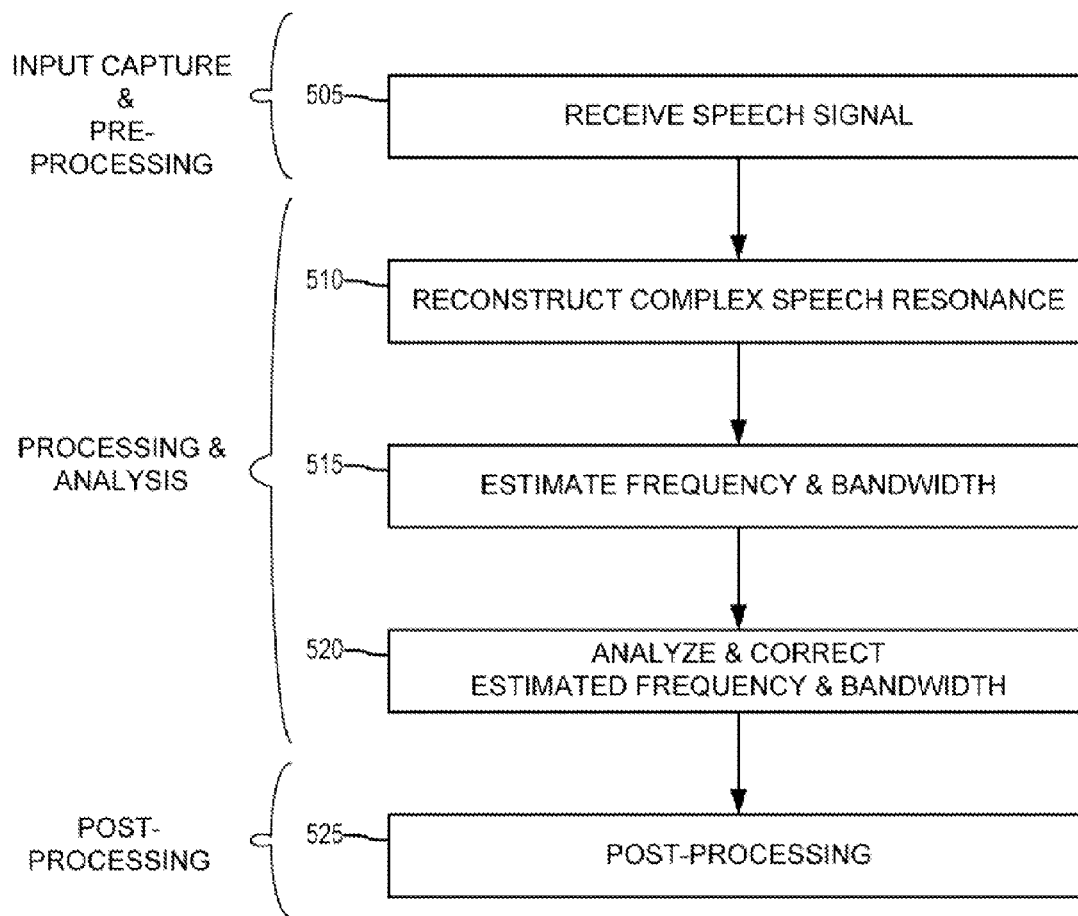
FIG. 5

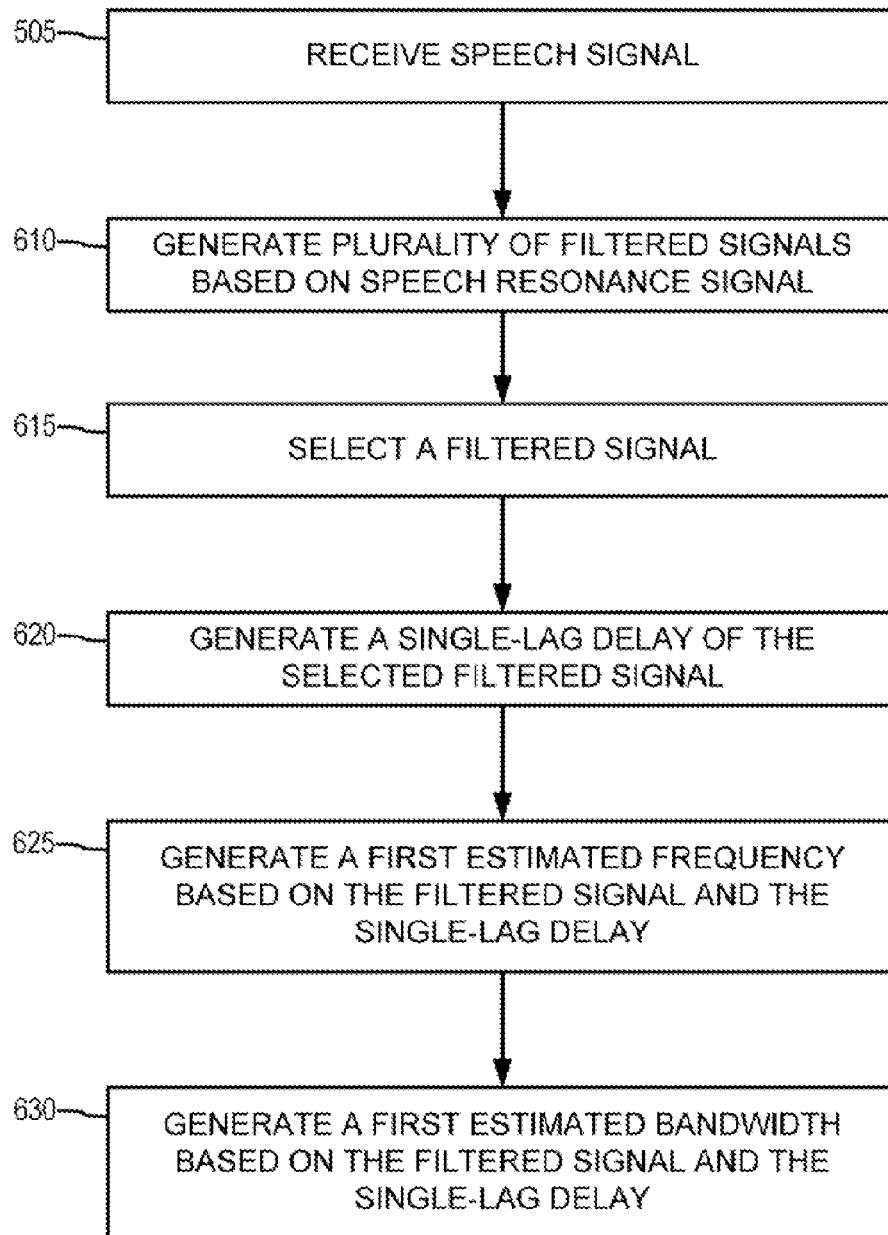
FIG. 6

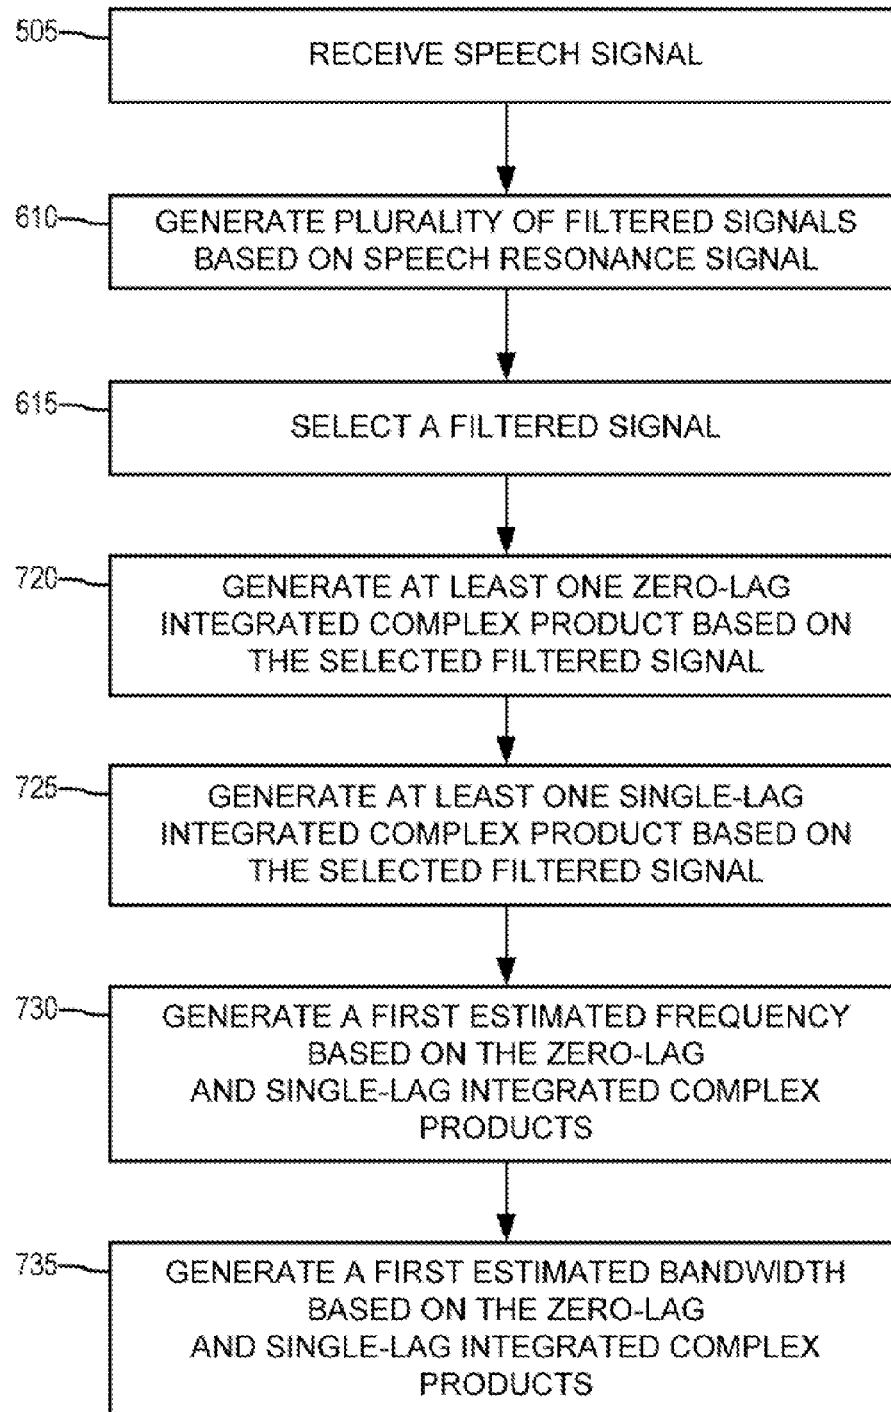
FIG. 7

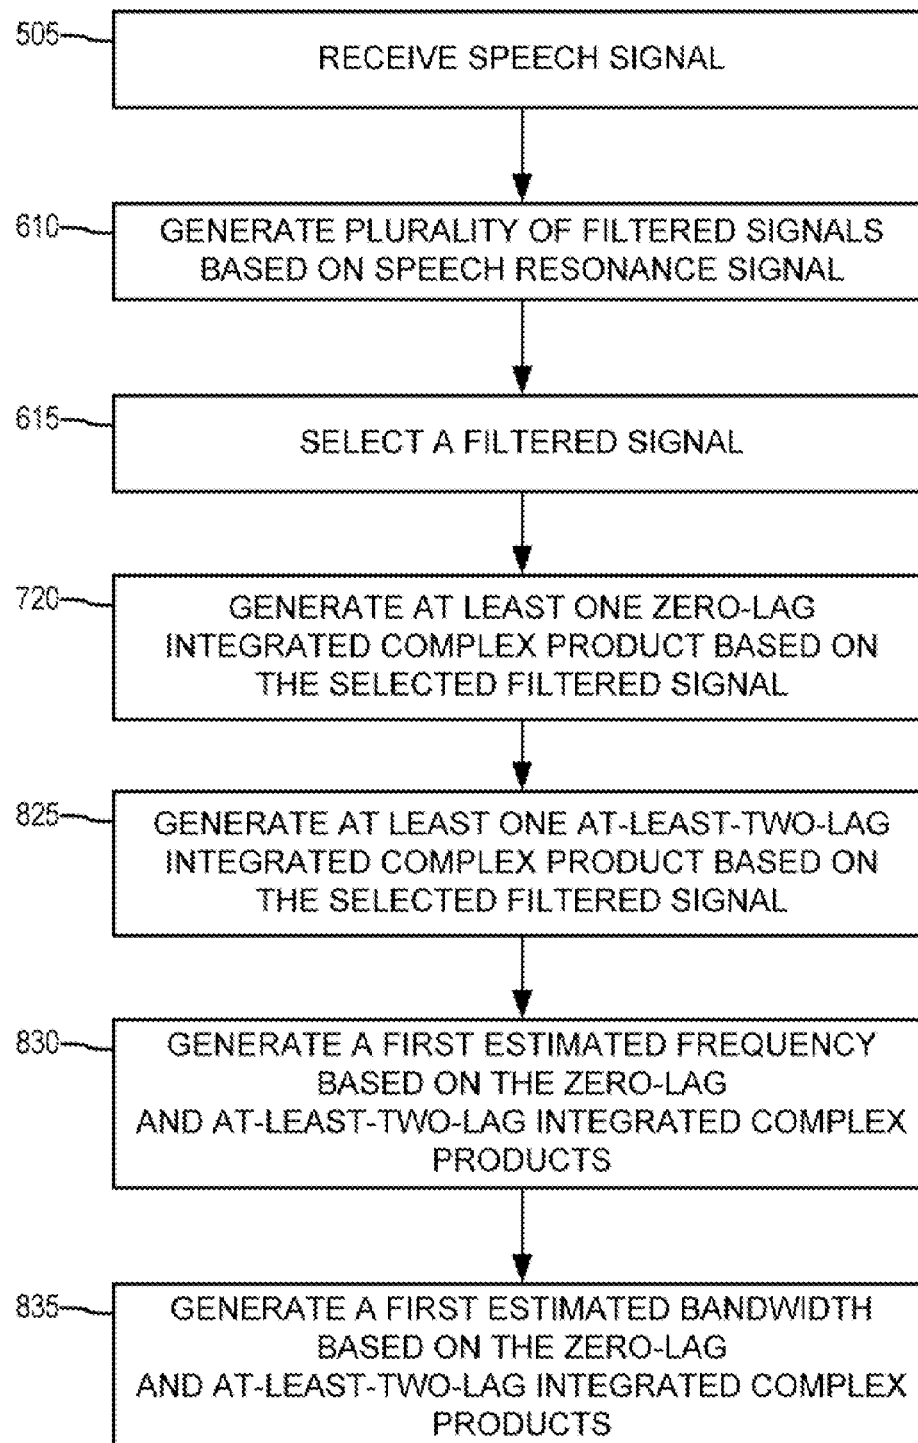
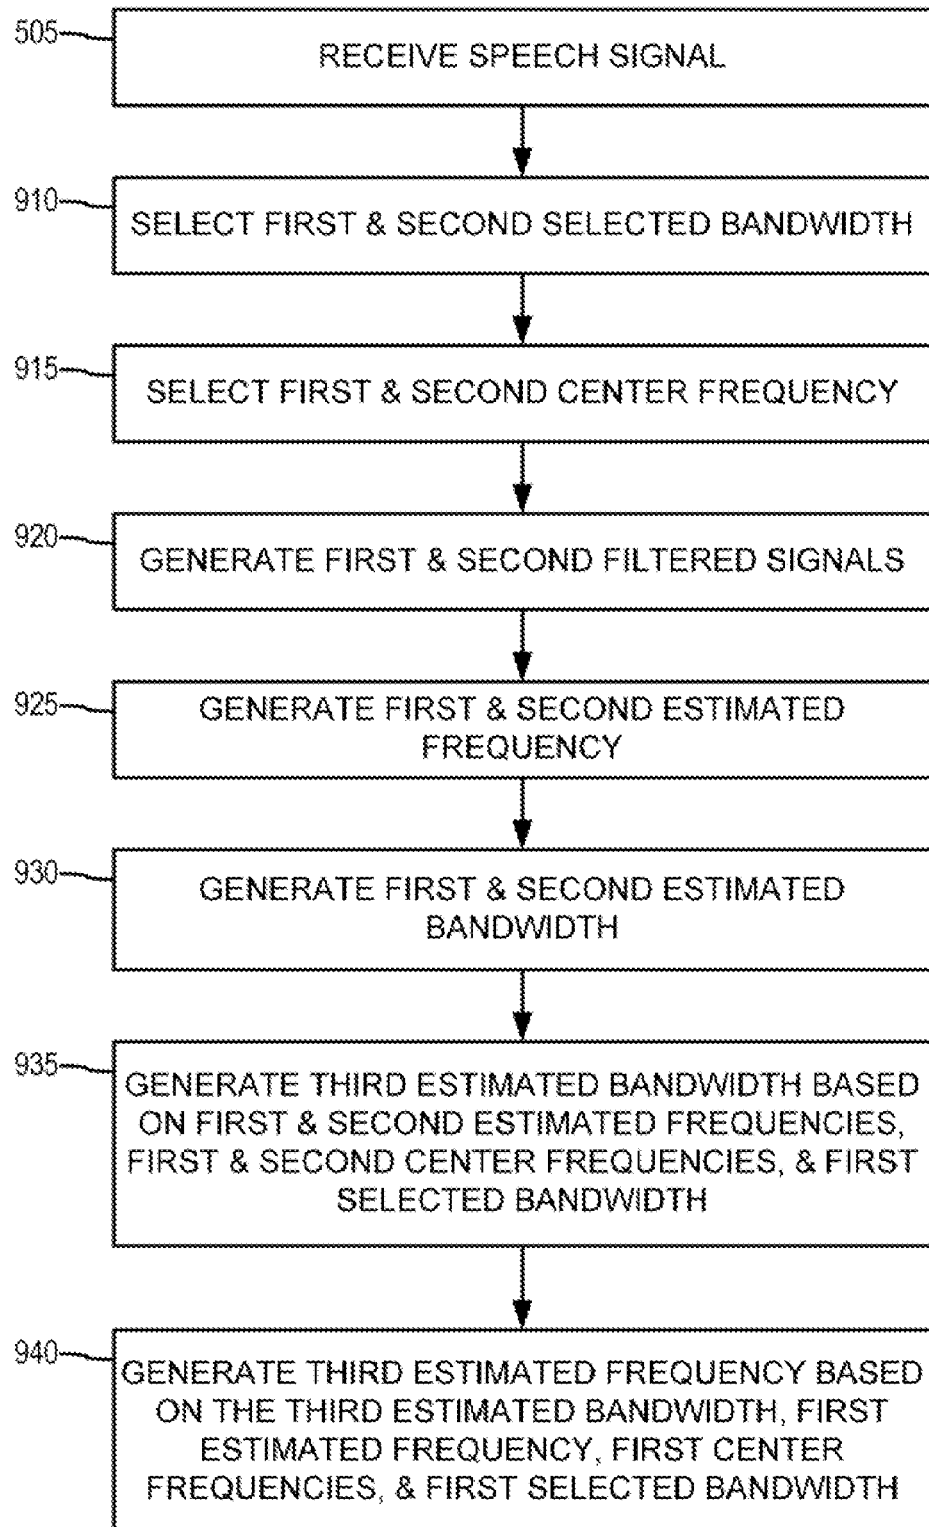
FIG. 8

FIG. 9

1

DIGITAL PROCESSOR BASED COMPLEX ACOUSTIC RESONANCE DIGITAL SPEECH ANALYSIS SYSTEM

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application is related to co-pending U.S. patent application Ser. No. 12/629,006, filed on Dec. 1, 2009 and which is incorporated herein by this reference.

FIELD OF THE INVENTION

The present invention relates generally to the field of speech recognition, and more particularly to systems for speech recognition signal processing and analysis.

BACKGROUND OF THE INVENTION

Modern human communication increasingly relies on the transmission of digital representations of acoustic speech over large distances. This digital representation contains only a fraction of the information about the human voice, and yet humans are perfectly capable of understanding a digital speech signal.

Some communication systems, such as automated telephone attendants and other interactive voice response systems (IVRs), rely on computers to understand a digital speech signal. Such systems recognize the sounds as well as the meaning inherent in human speech, thereby extracting the speech content of a digitized acoustic signal. In the medical and health care fields, correctly extracting speech content from a digitized acoustic signal can be a matter of life or death, making accurate signal analysis and interpretation particularly important.

One approach to analyzing a speech signal to extract speech content is based on modeling the acoustic properties of the vocal tract during speech production. Generally, during speech production, the configuration of the vocal tract determines an acoustic speech signal made up of a set of speech resonances. These speech resonances can be analyzed to extract speech content from the speech signal.

In order to determine accurately the acoustic properties of the vocal tract during speech production, both the frequency and the bandwidth of each speech resonance are required. Generally, the frequency corresponds to the size of the cavity within the vocal tract, and the bandwidth corresponds to the acoustic losses of the vocal tract. Together, these two parameters determine the formants of speech.

During speech production, speech resonance frequency and bandwidth may change quickly, on the order of a few milliseconds. In most cases, the speech content of a speech signal is a function of sequential speech resonances, so the changes in speech resonances must be captured and analyzed at least as quickly as they change. As such, accurate speech analysis requires simultaneous determination of both the frequency and bandwidth of each speech resonance on the same time scale as speech production, that is, on the order of a few milliseconds. However, the simultaneous determination of frequency and bandwidth of speech resonances on this time scale has proved difficult.

Some previous work in formant estimation has been concerned with finding only the frequency of speech resonances in speech signals. These frequency-oriented methods use the instantaneous frequency for high time-resolution frequency

2

estimates. However, these methods for frequency estimation are limited in flexibility, and do not fully describe the speech resonances.

For example, Nelson, et al., have developed a number of methods, including U.S. Pat. No. 6,577,968 for a "Method of estimating signal frequency," on Jun. 10, 2003, by Douglas J. Nelson; U.S. Pat. No. 7,457,756 for a "Method of generating time-frequency signal representation preserving phase information," on Nov. 25, 2008, by Douglas J. Nelson and David Charles Smith; and U.S. Pat. No. 7,492,814 for a "Method of removing noise and interference from signal using peak picking," on Feb. 17, 2009, by Douglas J. Nelson.

Generally, systems consistent with the Nelson methods ("Nelson-type systems") use instantaneous frequency to enhance the calculation of a Short-Time Fourier Transform (STFT), a common transform in speech processing. In Nelson-type systems, the instantaneous frequency is calculated as the time-derivative of the phase of a complex signal. The Nelson-type systems approach computes the instantaneous frequency from conjugate products of delayed whole spectra. Having computed the instantaneous frequency of each time-frequency element in the STFT, the Nelson-type systems approach re-maps the energy of each element to its instantaneous frequency. This Nelson-type re-mapping results in a concentrated STFT, with energy previously distributed across multiple frequency bands clustering around the same instantaneous frequency.

Auger & Flandrin also developed an approach, which is described in: F. Auger and P. Flandrin, "Improving the readability of time-frequency and time-scale representations by the reassignment method," *Signal Processing, IEEE Transactions on* 43, no. 5 (May 1995): 1068-1089 ("Auger/Flandrin"). Systems consistent with the Auger/Flandrin approach ("Auger/Flandrin-type systems") offer an alternative to the concentrated Short-Time Fourier Transform (STFT) of Nelson-type systems. Generally, Auger/Flandrin-type systems compute several STFTs with different windowing functions. Auger/Flandrin-type systems use the derivative of the window function in the STFT to get the time-derivative of the phase, and the conjugate product is normalized by the energy. Auger/Flandrin-type systems yield a more exact solution for the instantaneous frequency than the Nelson-type systems' approach, as the derivative is not estimated in the discrete implementation.

However, as extensions of STFT approaches, both Nelson-type and Auger/Flandrin-type systems lack the necessary flexibility to model human speech effectively. For example, the transforms of both Nelson-type and Auger/Flandrin-type systems determine window length and frequency spacing for the entire STFT, which limits the ability to optimize the filter bank for speech signals. Moreover, while both types find the instantaneous frequencies of signal components, neither type finds the instantaneous bandwidths of the signal components. As such, both the Nelson-type and Auger/Flandrin-type approaches suffer from significant drawbacks that limit their usefulness in speech processing.

Gardner and Mognasco describe an alternate approach in: T. J. Gardner and M. O. Mognasco, "Instantaneous frequency decomposition: An application to spectrally sparse sounds with fast frequency modulations," *The Journal of the Acoustical Society of America* 117, no. 5 (2005): 2896-2903 ("Gardner/Mognasco"). Systems consistent with the Gardner/Mognasco approach ("Gardner/Mognasco-type systems") use a highly-redundant complex filter bank, with the energy from each filter remapped to its instantaneous frequency, similar to the Nelson approach above. Gardner/Mo-

gnasco-type systems also use several other criteria to further enhance the frequency resolution of the representation.

That is, the Gardner/Mognasco-type systems discard filters with a center frequency far from the estimated instantaneous frequency, which can reduce the frequency estimation error from filters not centered on the signal component frequency. Gardner/Mognasco-type systems also use an amplitude threshold to remove low-energy frequency estimates and optimize the bandwidths of filters in a filter bank to maximize the consensus of the frequency estimates of adjacent filters. Gardner/Mognasco-type systems then use consensus as a measure of the quality of the analysis, where high consensus across filters indicates a good frequency estimate.

However, Gardner/Mognasco-type systems also suffer from significant drawbacks. First, Gardner/Mognasco-type systems do not account for instantaneous bandwidth calculation, thus missing an important part of the speech formant. Second, a consensus approach can lock in an error when a group of frequency estimates are briefly consistent with each other, but nevertheless provide inaccurate estimates of the true resonance frequency. For both of these reasons, Gardner/Mognasco-type systems offer limited usefulness in speech processing applications, particularly those applications that require higher accuracy over a short time scale.

While the above methods attempt to determine instantaneous frequency without also determining instantaneous bandwidth, Potamianos and Maragos developed a method for obtaining both the frequency and bandwidth of formants of a speech signal. The Potamianos/Maragos approach is described in: Alexandros Potamianos and Petros Maragos, "Speech formant frequency and bandwidth tracking using multiband energy demodulation," *The Journal of the Acoustical Society of America* 9, no. 6 (1996): 3795-3806 ("Potamianos/Maragos").

Systems consistent with the Potamianos/Maragos approach ("Potamianos/Maragos-type systems") use a filter bank of real-valued Gabor filters, and calculate the instantaneous frequency at each time-sample using an energy separation algorithm to demodulate the signal into an instantaneous frequency and amplitude envelope. In Potamianos/Maragos-type systems, the instantaneous frequency is then time-averaged to give a short-time estimate of the frequency, with a time window of about 10 ms. In Potamianos/Maragos-type systems, the bandwidth estimate is simply the standard deviation of the instantaneous frequency over the time window.

Thus, while Potamianos/Maragos-type systems offer the flexibility of a filter bank (rather than a transform), Potamianos/Maragos-type systems only indirectly estimate the instantaneous bandwidth by using the standard deviation. That is, because the standard deviation requires a time average, the bandwidth estimate in Potamianos/Maragos-type systems is not instantaneous. Because the bandwidth estimate is not instantaneous, the frequency and bandwidth estimates must be averaged over longer times than are practical for real-time speech recognition. As such, the Potamianos/Maragos-type systems also fail to determine speech formants on the time scale preferred for real-time speech processing.

SUMMARY OF THE INVENTION

In brief, the disclosed system extracts formants from a digital speech input signal by digitally filtering the speech signal substantially over its bandwidth to produce estimated instantaneous frequency and an instantaneous bandwidth information of resonances occurring in the speech signal in real time. Having received an analog speech signal, and hav-

ing sampled and digitized the samples, at least one digital processor is programmed to filter the speech signal using a plurality of computationally implemented complex digital filters to generate a plurality of complex digitally filtered signals. The bandwidths and center frequencies for each of the digital filters can be chosen such that they form a virtual chain of filters overlapping each other to ensure that substantially the entire relevant bandwidth of the of the speech signal is filtered by the chain. For each filtered digital signal, the at least one digital processor reconstructs a real component and an imaginary component of the speech signal. A single-lag delay of the speech signal is also generated, based on a selected filtered signal. The estimated frequency and bandwidth of speech resonances occurring in the speech signal are identified in real-time by the digital processor based on the estimated frequency and bandwidth of those resonances.

In one general aspect of the invention, a speech processing system extracts speech content from a digital speech signal. The speech content is characterized by at least one formant, and each of the at least one formants are characterized by an instantaneous frequency and an instantaneous bandwidth. The speech signal includes a sequence of one or more of the at least one formants. The speech processing system includes at least one digital processor. The at least one digital processor is programmed with instructions stored on at least one readable storage medium. The execution of the instructions by the at least one digital processor causes the digital processor to perform a method that includes extracting each one of the sequence of one or more of the at least one formants from the digital speech signal. The extracting process further includes filtering the digital speech signal using a plurality of complex digital filters, the plurality of digital filters being implemented to perform their digital filtering functions in parallel. Each of the digital filters has a predetermined bandwidth that covers an incremental portion of a total bandwidth of the digital speech signal. Each predetermined bandwidth overlaps with at least one other of the predetermined bandwidths. Each of the complex digital filters generates one of a plurality of complex digitally filtered signals. Each of the complex digitally filtered signals includes a real component and an imaginary component.

The extracting process further includes estimating an instantaneous frequency and an estimated instantaneous bandwidth from each of the plurality of digitally filtered signals using a product set formed of each of the plurality of digitally filtered signals in combination with a single lag delay of each of the plurality of digitally filtered signals. The extracting process further includes identifying each of the sequence of one or more formants of the digital speech signal as one of the at least one formants based on the estimated instantaneous frequencies and estimated instantaneous bandwidths. The system then reconstructs the speech content of the digital speech signal based on the identified sequence of formants.

In a further embodiment, the overlapping predetermined bandwidths of the plurality of complex digital filters taken together extend substantially over the bandwidth of the digital speech signal.

In another embodiment, at least one of the plurality of complex digital filters is characteristic of a finite impulse response (FIR) filter.

In another embodiment, at least one of the plurality of complex digital filters is characteristic of an infinite impulse response (IIR) filter.

In a further embodiment, at least one of the plurality of complex digital filters is characteristic of a gammatone filter.

5

In another aspect of the invention, the predetermined bandwidth of each of the complex digital filters is further characterized by a predetermined center frequency. The predetermined center frequency of each of the complex digital filters is separated by a predetermined center frequency spacing from the predetermined center frequency of the at least one of the plurality complex digital filters having a predetermined bandwidth that overlaps therewith. In one embodiment, the predetermined center frequency spacing is approximately 2%. In another embodiment, the predetermined bandwidth of each of the complex filters forming the chain is approximately 0.75 of its predetermined center frequency.

In one embodiment, the at least one digital processor is a general purpose microprocessor. In an alternate embodiment, the at least one digital processor is a digital signal processor (DSP) having computational resources designed to handle specific calculations intrinsic to said filtering and said estimating.

In a further embodiment, the generating process further includes integrating the product sets formed for each of the plurality of digitally filtered signals over a predetermined period of time to generate the estimated instantaneous frequency and the instantaneous bandwidth for each of digitally filtered signals.

In another embodiment, the generating further includes correcting the estimated instantaneous bandwidth for each one of the digitally filtered signals generated by one of the complex digital filters by first determining a difference between the estimated instantaneous frequency for two of the digitally filtered signals generated by digital filters having bandwidths overlapping the bandwidth of the one of the digital filters that generated the digitally filtered signal being corrected; secondly, by then dividing the determined difference by the predetermined center frequency spacing.

In another aspect of the invention, an integrated-product set is formed for each of the plurality of complex digitally filtered signals using an integration kernel, the integrated-product set having at least one zero-lag complex product and at least one single-lag complex product.

In still another embodiment, the integrated-product set has at least one zero-lag complex product and at least one two-or-more-lag complex product in place of the at least one single-lag complex product.

In yet another aspect of the invention, an apparatus extracts speech content embedded within a digitized speech signal, the speech content being characterized by at least one formant, each of the at least one formants characterized by an instantaneous frequency and an instantaneous bandwidth. The speech signal includes a sequence of one or more of the at least one formants. The apparatus includes a reconstruction processor configured by program instructions to receive and operate on samples of the digital speech signal. The reconstruction processor computationally implements a plurality of complex digital filters, the plurality of complex digital filters implemented to perform their processing in parallel on each sample of the digital speech signal. Each of the complex digital filters are characterized by a bandwidth that overlaps with the bandwidth of at least one other of the plurality of complex filters. Each of the complex digital filters generating as an output one of a plurality of digitally filtered signals. Each of the digitally filtered signals made up of discrete values for each sample of the digital speech signal processed, each of the digitally filtered signals including a real component and an imaginary component.

The apparatus further includes an estimator processor configured by program instructions to receive the plurality of digitally filtered signals from the reconstruction processor,

6

the estimator processor computationally implementing an estimator process, the estimator process being instantiated for each one of the generated digitally filtered signals, each instantiation of the estimator process configured to generate an estimated instantaneous frequency and an estimated instantaneous bandwidth from each of the plurality of digitally filtered signals using a product set formed of each of the plurality of digitally filtered signals.

The apparatus further includes a post-processing processor configured by program instructions to receive the estimated instantaneous frequency and instantaneous bandwidth estimates for each of the plurality of digitally filtered signals from the estimator processor. The post-processing processor further configured by program instructions to identify each of the sequence of one or more formants of the digital speech signal as one of the at least one formants based on the received estimated instantaneous frequencies and estimated instantaneous bandwidths of the plurality of filtered signals. The post-processing processor also configured by program instructions to reconstruct the speech content of the digital speech signal using the identified formants.

In an embodiment, each instantiation of the estimator process further comprises a computationally implemented integration kernel configured to integrate the product sets formed for each of the plurality of filtered signals over a predetermined period of time to generate the estimated instantaneous frequency and the instantaneous bandwidth for each of filtered signals.

In one embodiment, the integration kernel is characteristic of a second order gamma IIR filter.

In another embodiment, the estimated instantaneous frequency and the estimated instantaneous bandwidth from each of the plurality of digitally filtered signals is generated using a product set formed by the estimator process from each of the plurality of filtered signals in combination with at least one single lag-delay of each of the plurality of digitally filtered signals.

In a further embodiment, the estimator processor is further configured to implement a correction process that receives the estimated instantaneous frequency and the estimated instantaneous bandwidth from the estimator processor. The correction process provides a corrected estimated instantaneous bandwidth for each of the filtered signals to the post-processing module using a difference between the estimated instantaneous frequency for two adjacent complex filters in the chain divided by the predetermined center frequency spacing.

In still another embodiment, the correction process further provides a corrected estimated instantaneous frequency for each of the filtered signals to the post-processing processor by applying the corrected bandwidth for each of the filtered signals in a best-fit equation.

In another embodiment, the reconstruction processor, the estimator processor and the post-processing processor are implemented as one or more digital processors.

In an alternate embodiment, at least one of the one or more digital processors is a general purpose microprocessor.

In still another alternate embodiment, the reconstruction processor, the estimator processor and the post-processing processor are implemented as one or more DSP components.

BRIEF DESCRIPTION OF THE DRAWINGS

The embodiments described herein will be more fully understood by reference to the detailed description, in conjunction with the following figures, wherein:

FIG. 1a is a cutaway view of a human vocal tract;

7

FIG. 1*b* is a high-level block diagram of a speech processing system that includes a complex acoustic resonance speech analysis system;

FIG. 2 is a block diagram of an embodiment of the speech processing system of FIG. 1*b*, highlighting signal transformation and process organization;

FIG. 3*a* is a block diagram of an embodiment of a single digital processor based implementation of a speech resonance analysis process of the speech processing system of FIG. 2;

FIG. 3*b* is a block diagram of an embodiment of a distributed digital processor based implementation of a speech resonance analysis process of the speech processing system of FIG. 2;

FIG. 4 is a block diagram of an embodiment of a complex gammatone filter of a speech resonance analysis process;

FIG. 5 is a high-level flow diagram depicting operational steps of a speech processing method; and

FIGS. 6-9 are high-level flow diagrams depicting operational steps of embodiments of complex acoustic speech resonance analysis methods.

DETAILED DESCRIPTION

FIG. 1*a* illustrates a cutaway view of a human vocal tract 10. As shown, vocal tract 10 produces an acoustic wave 12. The qualities of acoustic wave 12 are determined by the configuration of vocal tract 10 during speech production. Specifically, as illustrated, vocal tract 10 includes four resonators 1, 2, 3, 4 that each contribute to generating acoustic wave 12. The four illustrated resonators are the pharyngeal resonator 1, the oral resonator 2, the labial resonator 3, and the nasal resonator 4. All four resonators, individually and together, create speech resonances during speech production. These speech resonances contribute to form the acoustic wave 12.

FIG. 1*b* illustrates an example of a speech processing system 100, in accordance with one embodiment of the invention. Broadly, speech processing system 100 operates in three general processing stages, "input capture and pre-processing," "processing and analysis," and "post-processing." Speech processing system 100 can be implemented using standard analog hardware components such as transistors, inductors, resistors and capacitors, one or more digital processors such as general purpose microprocessors (μ P) and/or application specific digital signal processors (DSP), or a combination of all of the foregoing. Each processing stage is described in additional detail below.

For analog implementations of the processing stages, the functions provided by the processing stages are performed by the components themselves on the signals as they pass through the hardware. For digital implementations, the processes are largely performed computationally on digital samples of the speech signal being analyzed. The computations are performed by one or more such processors based on program instructions that are stored on readable memory components separate from, or integrated within, the digital processors.

The difference between DSP and microprocessor components lies primarily in the type of dedicated resources that are available for performing the computations specific to the task at hand. General purpose microprocessors typically have generalized computational resources. DSP components tend to have computational resources that are more specifically tailored to performing the computations typically required for signal processing, and therefore tend to be faster but also more expensive. Both types of processing components are able to perform the computations necessary to the processing

8

stages as described herein, with general purpose processors tending to be slower and less expensive, and DSP components tending to be faster but more expensive. Thus, the use of the term digital processor hereinafter will be intended to cover any type of processing component capable of performing the computations requisite to the processing stages as described herein, including both general purpose microprocessors and application specific DSPs.

To analyze and interpret a speech signal, some speech must first be captured. The first stage of the process is therefore, generally, "input capture and pre-processing." As illustrated, speech processing system 100 is configured to capture acoustic wave 12, originating from vocal tract 10. As described above, a human vocal tract generates resonances in a variety of locations. In this stage, vocal tract 10 generates acoustic wave 12. Input processing module 110 detects, captures, and converts acoustic wave 12 into a digital speech signal.

More specifically, an otherwise conventional input processing module 110 captures the acoustic wave 12 through an input port 112. Input port 112 is an otherwise conventional input port and/or device, such as a conventional microphone or other suitable device. Input port 112 captures acoustic wave 12 and creates an analog signal 114 based on the acoustic wave.

Input processing module 110 also includes a digital distribution module 116. In one embodiment, digital distribution module 116 is an otherwise conventional device or system configured to digitize and distribute an input signal. Module 116 could be a separate or integrated analog-to-digital converter (ADC) as is known in the art. As shown, digital distribution module 116 receives analog signal 114 and generates an output signal 120 that consists of digitized samples of the analog signal 114, the samples typically being generated at a substantially constant sampling rate. In the illustrated embodiment, the output signal 120 is the output of input processing module 110.

The speech resonance analysis module 130 of the invention described herein receives the speech signal 120, forming an output signal suitable for additional speech processing by post processing module 140. As described in more detail below, speech resonance analysis module 130 reconstructs the speech signal 120 into a complex speech signal. Using the reconstructed complex speech signal, speech resonance analysis module 130 estimates the frequency and bandwidth of speech resonances of the complex speech signal, and can correct or further process the signal to enhance the accuracy of those estimates.

Speech resonance analysis module 130 passes its output to a post processing module 140, which can be configured to perform a wide variety of transformations, enhancements, and other post-processing functions, including the identification of formants within the output signal generated by speech resonance analysis module 130. In some embodiments, post processing module 140 is an otherwise conventional post-processing module. The following figures provide additional detail describing the invention.

FIG. 2 presents the processing and analysis stage in a representation capturing three broad processing sub-stages: reconstruction, estimation, and analysis/correction. Specifically, FIG. 2 shows another view of system 100. Input processing module 110 receives a real, analog, acoustic signal (i.e., a sound, speech, or other noise), captures the acoustic signal, converts it to a sampled digital format, and passes the resultant digital speech signal 120 to speech resonance analysis module 130.

One skilled in the art will understand that an acoustic resonance field such as human speech can be modeled as a

complex signal, and therefore can be described with a real component and an imaginary component. Generally, the input to input processing module 110 is a real, analog signal from, for example, the point 10 representing the vocal tract of FIG. 1, having lost the complex information during transmission. As shown, the output signal of module 110, speech signal 120 (shown as X), is a sampled digital representation of the analog input signal, and lacks some of the original signal information.

Speech signal 120 (signal X) is the input to the three stages of processing of the invention disclosed herein, referred to herein as “speech resonance analysis.” Specifically, reconstruction process 210 receives and reconstructs signal 120 such that the imaginary component and real components of each resonance are reconstructed. This stage is described in more detail below with respect to FIGS. 3a, 3b and 4. As shown, the output of reconstruction process 210 is a plurality of reconstructed digital signals Y_n , which each include a real component, Y_R , and an imaginary component, Y_I .

The output of the reconstruction process 210 is the input to the next broad stage of processing of the invention disclosed herein. Specifically, estimator process 210 receives signals Y_n , which is the output of the reconstruction stage. Very generally, estimator process 210 uses the reconstructed signals to estimate the instantaneous frequency and the instantaneous bandwidth of one or more of the individual speech resonances of the reconstructed speech signal. This stage is described in more detail below with respect to FIGS. 3a and 3b. As shown, the output of estimator process 210 is a plurality of estimated frequencies ($\hat{f}_1 \dots \hat{f}_n$) and estimated bandwidths ($\hat{\beta}_1 \dots \hat{\beta}_n$).

The output of the estimator process 210 is the input to the next broad stage of processing of the invention disclosed herein. Specifically, analysis & correction process 230 receives the plurality of estimated frequencies and bandwidths that are the output of the estimation stage. Very generally, process 230 uses the estimated frequencies and bandwidths to generate revised estimates. In one embodiment, the revised estimated frequencies and bandwidths are the result of novel corrective methods of the invention. In an alternate embodiment, the revised estimated frequencies and bandwidths, themselves the result of novel estimation and analysis methods, are passed to a post-processing module 140 for further refinement. This stage is described in more detail with respect to FIGS. 3a and 3b.

Generally, as described in more detail below, the output of the analysis and correction process 230 provides significant improvements over prior art systems and methods for estimating speech resonances. Configured in accordance with the invention described herein, a speech processing system can produce, and operate on, more accurate representations of human speech. Improved accuracy in capturing these formants results in better performance in speech applications relying on those representations.

More particularly, the invention presented herein determines individual speech resonances with a multi-object, parallel processing chain of digitally represented transfer functions that uses complex numbers throughout. Based on the properties of acoustic resonances, the invention is optimized to extract the frequency and bandwidth of speech resonances with high time-resolution.

FIGS. 3a and 3b illustrate embodiments of the invention in additional detail, which are implemented with digital processing components. In FIG. 3a, all of the speech analysis processes 130 (i.e. reconstruction process 210, estimator process 220, and analysis and correction process 230) are performed by a single digital processor 130p. In FIG. 3b, the

processing resources are more distributed so that each of the foregoing speech analysis processes are performed by a separate digital processor: reconstruction processor 210p, estimator processor 220p and analysis and correction processor 230. Those of skill in the art will appreciate that the distribution of such computational resources is primarily based on design considerations such as the speed with which computations must be made versus the cost of using multiple components to increase computational throughput.

Those of skill in the art will appreciate that a further embodiment can employ a separate processor for each of the computational processes represented by the complex digital filter functions 310 and each of the estimator processes 320 can be implemented as a separate processor. Another embodiment can implement each pairing of a complex digital filter function 310 and an estimator 320 together with a single digital processor.

Generally, speech recognition system 100 includes input processing module 110, which is configured to generate speech signal 120, as described above. As illustrated, reconstruction process 210 receives speech signal 120. In one embodiment, speech signal 120 is a digital speech signal sampled and digitized from a microphone or network source. In one embodiment, speech signal 120 is relatively low in accuracy and sampling frequency, e.g., 8-bit sampling. Reconstruction process 210 reconstructs the acoustic speech resonances using a general model of acoustic resonance.

For example, an acoustic resonance can be mathematically modeled as a complex exponential:

$$r(t) = e^{-2\pi\beta \cdot t} e^{-i2\pi f \cdot t}, \text{ for } t > 0$$

Where f is the frequency of the resonance (in Hertz), and β is the bandwidth (in Hertz). By convention, β is approximately the measurable full-width-at-half-maximum bandwidth. Further, complex sound transmission can be well described by a (real) sine wave. The signal capture process is thus the equivalent of taking the real (or imaginary) part of the complex source, which, however, also loses instantaneous information. As described in more detail below, reconstruction module 210 recreates the original complex representation of the acoustic speech resonances.

In the illustrated embodiment, reconstruction process 210 includes a plurality of complex digital filters (CFs) 310. Each of these complex filters are implemented digitally as a transfer function that characterizes the behavior of each filter, and each is applied computationally to each sample of the digital speech signal being processed simultaneously. One embodiment of a complex digital filter 310 is described in more detail with respect to FIG. 4, below. Generally, reconstruction process 210 produces a plurality of reconstructed signals, Y_n , each of which includes a real part (Y_R) and an imaginary part (Y_I).

As shown, system 100 includes an estimator process 220, which in the illustrated embodiment includes a plurality of estimator objects or instantiations 320, each of which is configured to receive sequential samples of one of the reconstructed signals Y_n . In the illustrated embodiment, each estimator object 320 includes an integration kernel 322. In an alternate embodiment, process 210 includes a single instantiation of estimator object 320, which can be configured with one or more integration kernels 322. In an alternate embodiment, estimator object 320 does not include an integration kernel 322. Those of skill in the art will appreciate that the computations performed by the estimator process can be performed in parallel by running n instantiations of the estimator process simultaneously. While the term “object” is used out of convenience to describe these separate instantiations of the

11

estimator process for each of the n signals, it is not intended that such processes must necessarily be the result of "object-oriented programming"

Generally, estimator objects **320** generate estimated instantaneous frequencies and bandwidths based on the reconstructed signals using the properties of an acoustic resonance. The equation for a complex acoustic resonance described above can be reduced to a very simple form:

$$r(t)=e^{-at}, \text{ with } a=2\pi\beta t+i2\pi f$$

for a resonance at frequency f , with bandwidth β . An equation of the family e^{-at} can also be modeled by a difference equation,

$$y[t]=(1-a)y[t-1]+x[t]$$

for a forcing function x . And if $x(t)$ is zero, as in a ringing response of the vocal tract resonances to an impulse from the glottal pulse, for example, in one embodiment, system **100** can determine the coefficient a based on two samples of a reconstructed resonance y , and from the coefficient a , the frequency and bandwidth can be estimated, as described in more detail below. In an alternate embodiment, also described in more detail below, where x is variable, or in noisy operating environment, system **100** can calculate auto-regression results to determine the coefficient a .

In the illustrated embodiment, each estimator object **320** passes the results of its frequency and bandwidth estimation to analysis and correction process **230**. Generally, process **230** receives a plurality of instantaneous frequency and bandwidth estimates and corrects these estimates, based on certain configurations, described in more detail below.

As shown, module **130** produces an output **340**, which, in one embodiment, system **100** sends to post processing module **140** for additional processing. In the embodiment, output **340** is a plurality of frequencies and bandwidths.

Thus, generally, system **100** receives a speech signal including a plurality of speech resonances, reconstructs the speech resonances, estimates their instantaneous frequency and bandwidth, and passes processed instantaneous frequency and bandwidth information on to a post-processing module for further processing, analysis, and interpretation. As described above, the first phase of analysis and processing is reconstruction, shown in more detail of one embodiment in FIG. 4.

FIG. 4 is a block diagram illustrating conceptual operation of a complex gammatone digital filter **310** in accordance with one embodiment. Specifically, filter **310** receives input speech signal **120**, divides speech signal **120** into two secondary input signals **412** and **414**, and passes the secondary input signals **412** and **414** through a series of filters **420**. In the illustrated embodiment, filter **310** includes a single series of filters **420**. In an alternate embodiment, filter **310** includes one or more additional series of filters **420**, arranged (as a series) in parallel to the illustrated series.

In the illustrated embodiment, the series of filters **420** is four filters long. So configured, the first filter **420** output serves as the input to the next filter **420**, which output serves as the input to the next filter **420**, and so forth.

In one embodiment, each filter **420** is a complex quadrature filter consisting of two filter sections **422** and **424**. In the illustrated embodiment, filter **420** is shown with two sections **422** and two sections **424**. In an alternate embodiment, filter **420** includes a single section **422** and a single section **424**, each configured to operate as described below. In one embodiment, each filter section **422** and **424** is a circuit configured to perform a transform on its input signal, described in more detail below. Each filter section **422** and **424** produces a

12

real number output, one of which applies to the real part of the filter **420** output, and the other of which applies to the imaginary part of the filter **420** output.

In one embodiment, filter **420** is a finite impulse response (FIR) filter. In one embodiment, filter **420** is an infinite impulse response (IIR) filter. In a preferred embodiment, the series of four filters **420** is a complex gammatone filter, which is a fourth-order gamma function envelope with a complex exponential. In an alternate embodiment, reconstruction module **310** is configured with other orders of the gamma function, corresponding to the number of filters **420** in the series.

Generally, the fourth-order gammatone filter impulse response is a function of the following terms:

$$g_n(t)=\text{Complex gammatone filter } n$$

$$b_n=\text{Bandwidth parameter of filter } n$$

$$f_n=\text{Center frequency of filter } n$$

and is given by:

$$g_n(t)=t^3 e^{-2\pi b_n t} e^{-i2\pi f_n t}, \text{ for } t>0$$

As such, in one embodiment, the output of filter **420** is an output of N complex numbers at the sampling frequency. Accordingly, the use of complex-valued filters eliminates the need to convert a real-valued input single into its analytic representation, because the response of a complex filter to a real signal is also complex. Thus, filter **310** provides a distinct processing advantage as filter **420** can be configured to unify the entire process in the complex domain.

Moreover, each filter **420** can be configured independently, with a number of configuration options, including the filter functions, filter window functions, filter center frequency, and filter bandwidth for each filter **420**. In one embodiment, the filter center frequency and/or filter bandwidth are selected from a predetermined range of frequencies and/or bandwidths. In one embodiment, each filter **420** is configured with the same functional form. In a preferred embodiment, each filter is configured as a fourth-order gamma envelope.

In one embodiment, each filter **420** filter bandwidth and filter spacing are configured to optimize overall analysis accuracy. As such, the ability to specify the filter window function, center frequency, and bandwidth of each filter individually contributes significant flexibility in optimizing filter **310**, particularly to analyze speech signals. In the preferred embodiment, each filter **420** is configured with 2% center frequency spacing and filter bandwidth of three-quarters of the center frequency (with saturation at 500 Hz). In one embodiment, filter **310** is a fourth-order complex gammatone filter, implemented as a cascade of first-order gammatone filters **420** in quadrature.

The following is a mathematical justification for using a cascade of first-order gammatone filters to create a fourth-order gammatone filter. For a complex input $x=x_R+ix_I$, the complex kernel of the first-order complex gammatone filter **420** can be represented as $g=g_R+ig_I$, where,

$$g_R(\tau)=e^{-2\pi b\tau} \cos 2\pi f\tau$$

$$g_I(\tau)=e^{-2\pi b\tau} \sin 2\pi f\tau$$

In one embodiment, filter sections **422** and **424** are configured respectively, with input signal s , as follows:

$$G_R(s)=\int g_R(\tau)s(t-\tau)d\tau$$

$$G_I(s)=\int g_I(\tau)s(t-\tau)d\tau$$

13

which, when combined, perform a first-order complex gammatone filter with output $y=y_R+iy_I$.

$$y_R(t)=G_R(x_R)-G_I(x_I)$$

$$y_I(t)=G_I(x_R)+G_R(x_I)$$

As such, in one embodiment, a fourth-order complex gammatone filter is four iterations of the first-order filter 420:

$$G_4(x)=G_1 \cdot G_1 \cdot G_1 \cdot G_1(x) \quad (4.4)$$

In the illustrated embodiment, for example, each filter 420 is configured as a first order gammatone filter. Specifically, filter 310 receives an input signal 120, and splits the received signal into designated real and imaginary signals. In the illustrated embodiment, splitter 410 splits signal 120 into a real signal 412 and an imaginary signal 414. In an alternate embodiment, splitter 410 is omitted and filter 420 operates on signal 120 directly. In the illustrated embodiment, both real signal 412 and “imaginary” signal 414 are real-valued signals, representing the complex components of input signal 120.

In the illustrated embodiment, real signal 412 is the input signal to a real filter section 422 and an imaginary filter 424. In the illustrated embodiment, section 422 calculates G_R from signal 412 and section 424 calculates G_I from signal 412. Similarly, imaginary signal 414 is the input signal to a real filter section 422 and an imaginary filter section 424. In the illustrated embodiment, section 422 calculates G_R from signal 414 and section 424 calculates G_I from signal 414.

As shown, filter 420 combines the outputs from sections 422 and 424. Specifically, filter 420 includes a signal subtractor 430 and a signal adder 432. In the illustrated embodiment, subtractor 430 and adder 432 are configured to subtract or add the signal outputs from sections 422 and 424. One skilled in the art will understand that there are a variety of mechanisms suitable for adding and/or subtracting two signals. As shown, subtractor 430 is configured to subtract the output of imaginary filter section 424 (to which signal 414 is input) from the output of real filter section 422 (to which signal 412 is input). The output of subtractor 430 is the real component, Y_R , of the filter 420 output.

Similarly, adder 432 is configured to add the output of imaginary filter section 424 (to which signal 412 is input) to the output of real filter section 422 (to which signal 414 is input). The output of adder 432 is the real value of the imaginary component, Y_I , of the filter 420 output. As shown, module 400 includes four filters 420, the output of which is a real component 440 and an imaginary component 442. As described above, real component 440 and imaginary component 442 are passed to an estimator module for further processing and analysis.

It will be appreciated by those of skill in the art that the foregoing filter implementations are realized as a computational process that is executed by a digital processor to generate the outputs of the complex digital filters 310, and that each instantiation of that computational process has its own bandwidth and center frequency such that the bandwidths of the plurality can be made to overlap with one another to ensure coverage of the entire bandwidth of the digital speech signal to be analyzed. By overlapping the bandwidths of adjacent instantiations of the digital filter 310 as a virtual chain, no resonance information contained within the input speech signal will pass through without being detected.

Returning now to FIGS. 3a and 3b, in the illustrated embodiment of system 100, estimator process 210 includes a plurality of estimator objects or instantiations 320. As described above, each estimator object 320 receives a real

14

component (Y_R) and a (real-valued) imaginary component (Y_I) from one of the complex digital filters 310 of reconstruction module 210. In one embodiment, each estimator object 320 receives or is otherwise aware of the configuration of the particular complex digital filter 310 that generated the input to that estimator object 320. In one embodiment, each estimator object 320 is associated with a complex filter 310, and is aware of the configuration setting of the complex filter 310, including the filter function(s), filter center frequency, and filter bandwidth.

In the illustrated embodiment, each estimator object 320 also includes an integration kernel 322, which adds an additional computational process to that performed by each estimator object 320. In an alternate embodiment, each estimator object 320 operates without an integration kernel 322. In one embodiment, at least one integration kernel 322 is a second order gamma IIR filter. Generally, each integration kernel 322 is configured to receive real and imaginary components as inputs, and to calculate zero-lag delays and variable-lag delays based on the received inputs.

Each estimator object 320 uses variable-delays of the filtered signals to form a set of products to estimate the frequency and bandwidth using methods described below. There are several embodiments of the estimator object 320; for example, the estimator object 320 may contain an integration kernel 322, as illustrated. For clarity, three alternative embodiments of the system with increasing levels of complexity are introduced here.

In the first embodiment, each estimator object 320 generates an estimated frequency and an estimated bandwidth of a speech resonance of the input speech signal 120 without an integration kernel 322. The estimated frequency and bandwidth are based only on the current filtered signal output from the CF 310 associated with that estimator object 320, and a single-lag delay of that filtered signal output. In one embodiment, the plurality of filters 310 and associated estimator objects 320 generate a plurality of estimated frequencies and bandwidths at each time sample.

In a second embodiment, each estimator object 320 includes an integration kernel 322, which forms an integrated-product set. Based on the integrated-product set, estimator object 320 generates an estimated frequency and an estimated bandwidth of a speech resonance of the input speech signal 120. Each integration kernel 322 forms the integrated-product set by updating products of the filtered signal output and a single-delay of the filtered signal output for the length of the integration. In one embodiment, the plurality of filters 310 and associated estimator objects 320 generate a plurality of estimated frequencies and bandwidths at each time sample, which are smoothed over time by the integration kernel 322.

In a third embodiment, the integrated-product set has an at-least-two-lag complex product, increasing the number of products in the integrated-product set. These three embodiments are described in more detail below.

In the first embodiment introduced above, estimator object 320 computes a single-lag product set using the output of a CF 310 without integration kernel 322. In this embodiment, the product set $\{y[t]y^*[t-1], |y[t]|^2\}$, where y is the complex output of CF 310, is used to find the instantaneous frequency and bandwidth of the input speech signal 102 using a single delay, extracting a single resonance at each point in time. Estimator object 320 computes the instantaneous frequency \hat{f} and instantaneous bandwidth $\hat{\beta}$ with the single-lag product set using the following equations:

15

$$\hat{f} = 2\pi dt \cdot \arg\left(\frac{y[t]y^*[t-1]}{|y[t]|^2}\right)$$

$$\hat{\beta} = -\frac{1}{2\pi dt} \ln\left(\frac{y[t]y^*[t-1]}{|y[t]|^2}\right)$$

where dt is the sampling interval. In a preferred embodiment, one or more estimator objects **320** calculate the instantaneous frequency and bandwidth from a single-lag product set based on each CF **310** output.

In alternate embodiments (e.g., the second and third embodiments introduced above), estimator object **320** computes an integrated-product set of variable delays using integration kernel **322**. The integrated-product set is used to compute the instantaneous frequency and bandwidth of the speech resonances of the input speech signal **102**. In a preferred embodiment, one or more estimator objects **320** calculate an integrated-product set based on each CF **310** output.

The integrated-product set of the estimator object **320** can include zero-lag products, single-lag products, and at-least-two lag products depending on the embodiment. In these embodiments, the integrated-product set is configured as an integrated-product matrix with the following definitions:

$\Phi_N(t)$ =Integrated-product matrix with N delays

$\Phi_{m,n}(t)$ =Integrated-product matrix element with delays $m, n \leq N$

y =Complex-signal output of CF **310** in reconstruction Module **210**

k =Integration kernel **322** within Estimator module **320**

Estimator object **320** updates the elements of the integrated-product matrix at each sampling time, with time-integration performed separately for each element over a integration kernel $k[\tau]$ of length l ,

$$\Phi_{m,n}(t) = \sum_{\tau=0}^l k[\tau] y[t-\tau-m] y^*[t-\tau-n]$$

The full integrated-product set with N -delays is an $N+1$ -by- $N+1$ matrix:

$$\Phi_N = \begin{bmatrix} \varphi_{0,0} & \dots & \varphi_{0,N} \\ & \dots & \\ \varphi_{N,0} & \dots & \varphi_{N,N} \end{bmatrix}$$

As such, for a maximum delay of 1 (i.e. a single-lag), the integrated product set is a 2×2 matrix:

$$\Phi_1 = \begin{bmatrix} \varphi_{0,0} & \varphi_{0,1} \\ \varphi_{1,0} & \varphi_{1,1} \end{bmatrix}$$

Accordingly, element $\varphi_{0,0}$ is a zero-lag complex product and elements $\varphi_{0,1}$, $\varphi_{1,1}$, and, $\varphi_{1,0}$ are single-lag complex products. Additionally, for a maximum delay of 2 (i.e., an at-least-two-lag), the integrated-product set is a 3×3 matrix, composed of the zero-lag and single-lag products from above, as well as an additional column and row of two-lag products: $\varphi_{0,2}$, $\varphi_{1,2}$, $\varphi_{2,2}$, $\varphi_{2,1}$, and, $\varphi_{2,0}$. Generally, additional lags

16

improve the precision of subsequent frequency and bandwidth estimates. One skilled in the art will understand that there is a computational trade-off between precision gained by additional lags and the power/time required to compute the additional elements.

In this embodiment, estimator object **320** is configured to use time-integration to calculate the integrated-product set. Generally, complex time-integration provides flexible optimization for estimates of speech resonances. For example, time-integration can be used to average resonance estimates over the glottal period to obtain more accurate resonance values, independent of glottal forcing.

Function k is chosen to optimize the signal-to-noise ratio while preserving speed of response. In a preferred embodiment, the integration kernel **322** configures k as a second-order gamma function. In one embodiment, integration kernel **322** is a second-order gamma IIR filter. In an alternate embodiment, integration kernel **322** is an otherwise conventional FIR or IIR filter.

In the second embodiment with a single-delay integrated-product set, introduced above, the estimator object **320** calculates the instantaneous frequency \hat{f} and instantaneous bandwidth $\hat{\beta}$ using elements of the single-delay integrated-product matrix with the following equations:

$$\hat{f} = 2\pi dt \cdot \arg\left(\frac{\varphi_{1,0}}{\varphi_{1,1}}\right)$$

$$\hat{\beta} = -\frac{1}{2\pi dt} \ln\left(\frac{\varphi_{1,0}}{\varphi_{1,1}}\right)$$

In this embodiment, $\hat{\beta}$ is the estimated bandwidth associated with a pole-model of a resonance. One skilled in the art will understand that other models can also be employed.

It is worth nothing that these equations for frequency and bandwidth estimation are equivalent to the equations in the first embodiment described above, where the integration window k is configured as a Kronecker delta function, essentially removing the integration kernel, resulting in the equivalent product matrix elements:

$$\Phi_{m,n}(t) = y[t-m] y^*[t-n]$$

In the third embodiment introduced above, estimator object **320** uses an integrated product-set with additional delays to estimate the properties of more resonances per complex filter at each sample time. This can be used in detecting closely-spaced resonances.

In summary, reconstruction module **310** provides an approximate complex reconstruction of an acoustic speech signal. Estimator objects **320** use the reconstructed signals that are the output of module **310** to compute the instantaneous frequency and bandwidth of the resonance, based in part on the properties of acoustic resonance generally.

In the illustrated embodiment, analysis and correction module **330** receives the plurality of estimated frequencies and bandwidths, as well as the product sets from the estimator objects **320**. Generally, analysis & correction module **330** provides an error estimate of the frequency and bandwidth calculations using regression analysis. The analysis & correction module uses the properties of the filters in recognition module **310** to produce one or more corrected frequency and bandwidth estimates **340** for further processing, analysis, and interpretation.

In one embodiment, analysis & correction module **230** processes the output of the integrated-product set as a complex auto-regression problem. That is, module **330** computes

17

the best difference equation model of the complex acoustic resonance, adding a statistical measure of fit. More particularly, in one embodiment, analysis & correction module **330** calculates an error estimate from the estimation objects **320** using the properties of regression analysis in the complex domain with the following equation:

$$r^2 = \frac{\varphi_{0,0} - \varphi_{1,1} \cdot \left| \frac{\varphi_{1,0}}{\varphi_{1,1}} \right|^2}{\varphi_{0,0}}$$

The error r is a measure of the goodness-of-fit of the frequency estimate. In one embodiment, module **330** uses r to identify instantaneous frequencies resulting from noise versus those resulting from resonance. Use of this information in increasing the accuracy of the estimates is discussed below.

In addition to an error estimate, an embodiment of analysis & correction module **230** also estimates a corrected instantaneous bandwidth of a resonance by using the estimates from one or more estimator objects **320**. In a preferred embodiment, module **230** estimates the corrected instantaneous bandwidth using pairs of frequency estimates, as determined by estimator objects **320** with corresponding complex filters **312** closely spaced in center frequency. Generally, this estimate better approximates the bandwidth of the resonance than the single-filter-based estimates described above.

Specifically, module **230** can be configured to calculate a more accurate bandwidth estimate using the difference in frequency estimate over the change in center frequency across two adjacent estimator modules,

$$v_n = \frac{\hat{f}_{n+1} - \hat{f}_n}{f_{n+1} - f_n}$$

The corrected instantaneous bandwidth estimate from the n^{th} estimator object **320**, $\hat{\beta}_n$, can be estimated using the selected bandwidth of the corresponding complex filter **312**, b_n , with the following equation:

$$\hat{\beta}_n = a_0 v_n \left(\frac{1 + a_1 v_n - a_2 v_n^2}{1 + a_3 v_n - a_4 v_n^2} \right) b_n$$

where, in one embodiment, the preferred coefficients, found empirically, are:

$$a_0 = 6.68002$$

$$a_1 = 3.69377$$

$$a_2 = 2.87388$$

$$a_3 = 47.5236$$

$$a_4 = 42.4272$$

In one embodiment, in particular where each CF **310** is a complex gammatone filter, the estimated instantaneous frequency can be skewed away from the exact value of the original resonance, in part because of the asymmetric frequency response of the complex filters **310**. Thus, module **230** can be configured to use the corrected bandwidth estimate, obtained using procedures described above, to correct errors in the estimated instantaneous frequencies coming from the

18

estimator objects **320**. For example, in one embodiment, for a CF **310** with center frequency f , bandwidth b , and uncorrected frequency estimate \hat{f} , the best-fit equation for frequency estimate correction is:

$$\hat{f}_{corrected} = f + (1 + 3.92524 \cdot R^2) \cdot (\hat{f} - f - c_1 R^2 \cdot e^{-c_3 R})$$

where $R = \hat{\beta}/b$ is the ratio of estimated resonance bandwidth to filter bandwidth. In one embodiment, the constants are found empirically. For example, where $b < 500$:

$$c_1 = 0.059101 + 0.816002 \cdot f$$

$$c_2 = 2.3357$$

$$c_3 = 3.58372$$

and for $b = 500$:

$$c_1 = 0.513951 + 140340.0 / (-409.325 + f)$$

$$c_2 = 1.95121 + 145.771 / (-292.151 + f)$$

$$c_3 = 1.72734 + 654.08 / (-319.262 + f)$$

As such, analysis and correction process **230** can be configured to improve the accuracy of the estimated resonance frequency and bandwidth generated by the estimator objects **320**. Thus, the improved estimates can be forwarded for speech recognition processing and interpretation, with improved results over estimates generated by prior art approaches.

For example, in one embodiment, post-processing module **140** performs thresholding operations on the plurality of estimates received from analysis & correction modules **230**. In one embodiment, thresholding operations discard estimates outside a predetermined range in order to improve signal-to-noise performance. In one embodiment, module **140** aggregates the received estimates to reduce the over-determined data-set. One skilled in the art will understand that module **140** can be configured to employ other suitable post-processing operations.

Thus, generally, system **100** can be configured to perform all three stages of speech signal process and analysis described above, namely, reconstruction, estimation, and analysis/correction. The following flow diagrams describe these stages in additional detail. Referring now to FIG. 5, the illustrated process begins at block **505**, in an input capture and pre-processing stage, wherein the speech recognition system receives a speech signal. For example, reconstruction process **210** receives a speech signal from input processing module **110** (of FIG. 2).

Next, the process enters the processing and analysis stage. Specifically, as indicated at block **510**, reconstruction process **210** reconstructs the received speech signal. Next, as indicated at block **515**, estimator process **210** estimates the frequency and bandwidth of a speech resonance of the reconstructed speech signal. Next, as indicated at block **520**, analysis and correction process **230** performs analysis and correction operations on the estimated frequency and bandwidth of the speech resonance.

Next, the process enters the post-processing stage. Specifically, as indicated at block **525**, post-processing module **140** performs post-processing on the corrected frequency and bandwidth of the speech resonance. Particular embodiments of this process are described in more detail below.

Referring now to FIG. 6, the illustrated process begins at block **505**, as above. Next, as indicated at block **610**, reconstruction process **210** generates a plurality of filtered signals based on a speech resonance signal of the received speech signal received as described in block **505**. In the preferred

19

embodiment, each of the plurality of filtered signal is a reconstructed (real and complex) speech signal, as described above.

Next, as indicated at block **615**, estimator process **210** selects one of the filtered signals generated as described in block **610**. Next, as indicated at block **620**, estimator process **210** generates a single-lag delay of a speech resonance of the selected filtered signal.

Next, as indicated at block **625**, estimator process **210** generates a first estimated frequency of the speech resonance based on the filtered signal and the single-lag delay of the selected filtered signal. Next, as indicated at block **630**, estimator process **210** generates a first estimated bandwidth of the speech resonance based on the filtered signal and the single-lag delay of the selected filtered signal. Thus, the flow diagram of FIG. 6 describes a process that generates an estimated frequency and bandwidth of a speech resonance of a speech signal.

Referring now to FIG. 7, the illustrated process advances as described above as indicated in blocks **505**, **610**, and **615**. Next, as indicated at block **720**, estimator process **210** generates at least one zero-lag integrated complex product based on the filtered signal selected as described in block **615**. Next, as indicated at block **725**, estimator process **210** generates at least one single-lag integrated complex product based on the selected filtered signal.

Next, as indicated at block **730**, estimator process **210** generates a first estimated frequency based on the zero-lag and single-lag integrated complex products. Next, as indicated at block **735**, estimator process **210** generates a first estimated bandwidth based on the zero-lag and single-lag integrated complex products.

Referring now to FIG. 8, the illustrated process advances as described above as indicated in blocks **505**, **610**, **615**, and **720**. Next, as indicated at block **825**, estimator process **210** generates at least one at-least-two-lag integrated complex product based on the selected filtered signal.

Next, as indicated at block **830**, estimator process **210** generates a first estimated frequency based on the zero-lag and at-least-two-lag integrated complex products. Next, as indicated at block **835**, estimator process **210** generates a first estimated bandwidth based on the zero-lag and at-least-two-lag integrated complex products.

Referring now to FIG. 9, the illustrated process begins as described above as indicated in block **505**. Next, as indicated at block **910**, reconstruction process **210** selects a first and second bandwidth. As described above, in one embodiment, reconstruction process **210** selects a first bandwidth, used to configure a first complex filter, and a second bandwidth, used to configure a second complex filter.

Next, as indicated at block **915**, reconstruction process **210** selects a first and second center frequency. As described above, in one embodiment, reconstruction process **210** selects a first center frequency, used to configure the first complex filter, and a second center frequency, used to configure the second complex filter. Next, as indicated at block **920**, reconstruction process **210** generates a first and second filtered signal. As described above, in one embodiment, the first filter generates the first filtered signal and the second filter generates the second filtered signal.

Next, as indicated at block **925**, estimator process **210** generates a first and second estimated frequency. As described above, in one embodiment, estimator process **210** generates a first estimated frequency based on the first filtered signal, and generates a second estimated frequency based on the second filtered signal.

20

Next, as indicated at block **930**, estimator process **210** generates a first and second estimated bandwidth. As described above, in one embodiment, estimator process **210** generates a first estimated bandwidth based on the first filtered signal, and generates a second estimated bandwidth based on the second filtered signal.

Next, as indicated at block **935**, analysis and correction process **230** generates a third estimated bandwidth based on the first and second estimated frequencies, the first and second center frequencies, and the first selected bandwidth. Next, as indicated at block **940**, analysis and correction process **230** generates a third estimated frequency based on the third estimated bandwidth, the first estimated frequency, the first center frequency, and the first selected bandwidth.

Other modifications and implementations will occur to those skilled in the art without departing from the spirit and scope of the invention as claimed. Accordingly, the above description is not intended to limit the invention except as indicated in the following claims.

What is claimed is:

1. A speech processing system for extracting speech content from a digital speech signal, the speech content being characterized by at least one formant, each of the at least one formants characterized by an instantaneous frequency and an instantaneous bandwidth, the speech signal including a sequence of one or more of the at least one formants, the speech processing system comprising:

at least one digital processor, the at least one digital processor programmed with instructions stored on at least one readable storage medium, the execution of the instructions by the at least one digital processor causing the at least one digital processor to perform the method of:

extracting each one of the sequence of one or more of the at least one formants from the digital speech signal, said extracting further comprising:

filtering the digital speech signal using a plurality of complex digital filters, the plurality of digital filters implemented to perform their digital filtering functions in parallel, each of the digital filters having a predetermined bandwidth that covers an incremental portion of a total bandwidth of the digital speech signal, each predetermined bandwidth overlapping with at least one other of the predetermined bandwidths, each of the complex digital filters generating one of a plurality of complex digitally filtered signals, each of the complex digitally filtered signals including a real component and an imaginary component;

generating an estimated instantaneous frequency and an estimated instantaneous bandwidth from each of the plurality of digitally filtered signals using a product set formed of each of the plurality of digitally filtered signals in combination with a single lag delay of each of the plurality of digitally filtered signals; and

identifying each of the sequence of one or more formants of the digital speech signal as one of the at least one formants based on the estimated instantaneous frequencies and estimated instantaneous bandwidths; and

reconstructing the speech content of the digital speech signal based on the identified sequence of formants.

2. The speech processing system of claim 1, wherein the overlapping predetermined bandwidths of the plurality of complex digital filters taken together extend substantially over the bandwidth of the digital speech signal.

21

3. The digital speech processing system of claim 1, wherein at least one of the plurality of complex digital filters is characteristic of a finite impulse response (FIR) filter.

4. The speech processing system of claim 1, wherein at least one of the plurality of complex digital filters is characteristic of an infinite impulse response (IIR) filter.

5. The speech processing system of claim 1, wherein at least one of the plurality of complex digital filters is characteristic of a gammatone filter.

6. The speech processing system of claim 1, wherein the predetermined bandwidth of each of the complex digital filters is further characterized by a predetermined center frequency, the predetermined center frequency of each of the complex digital filters being separated by a predetermined center frequency spacing from the predetermined center frequency of the at least one of the plurality complex digital filters having a predetermined bandwidth that overlaps therewith.

7. The speech processing system of claim 6, wherein the predetermined center frequency spacing is approximately 2%.

8. The speech processing system of claim 7, wherein the predetermined bandwidth of each of the plurality of complex filters is approximately 0.75 of its predetermined center frequency.

9. The speech processing system of claim 6 wherein said generating further comprises correcting the estimated instantaneous bandwidth for each one of the digitally filtered signals generated by one of the complex digital filters, said correcting further comprising:

determining a difference between the estimated instantaneous frequency for two of the digitally filtered signals generated by digital filters having bandwidths overlapping the bandwidth of the one of the digital filters that generated the digitally filtered signal being corrected; and

dividing the determined difference by the predetermined center frequency spacing.

10. The speech processing system of claim 1, wherein the at least one digital processor is a general purpose microprocessor.

11. The speech processing system of claim 1, wherein the at least one digital processor is a digital signal processor (DSP) having computational resources designed to handle specific calculations intrinsic to said filtering and said estimating.

12. The speech processing system of claim 1 wherein said generating further comprises integrating the product sets formed for each of the plurality of digitally filtered signals over a predetermined period of time to generate the estimated instantaneous frequency and the instantaneous bandwidth for each of digitally filtered signals.

13. A speech processing system for extracting speech content from a digital speech signal, the speech content being characterized by at least one formant, each of the at least one formants characterized by an instantaneous frequency and an instantaneous bandwidth, the speech signal including a sequence of one or more of the at least one formants, the system comprising:

at least one digital processor, the at least one digital processor programmed with instructions stored on at least one readable storage medium, the execution of the instructions by the at least one digital processor causing the at least one digital processor to perform the method of:

22

extracting each one of the sequence of formants from the digital speech signal, said extracting further comprising:

filtering the speech resonance signal with a plurality of complex digital filters, implemented with overlapping bandwidths to form a virtual parallel processing chain, to generate a plurality of complex digitally filtered signals having a real component and an imaginary component;

forming an integrated-product set for each of the plurality of complex digitally filtered signals using an integration kernel, the integrated-product set having at least one zero-lag complex product and at least one single-lag complex product;

generating an estimated instantaneous frequency and an estimated instantaneous bandwidth from each of the integrated-product sets; and

identifying each of the sequence of one or more formants of the digital speech signal as one of the at least one formants based on the estimated instantaneous frequencies and estimated instantaneous bandwidths; and

reconstructing the speech content of the digital speech signal based on the identified sequence of formants.

14. The speech processing system of claim 13, wherein at least one of the plurality of complex digital filters of the virtual parallel processing chain is characteristic of a finite impulse response (FIR) filter.

15. The speech processing system of claim 13, wherein at least one of the plurality of complex digital filters of the virtual parallel processing chain is characteristic of an infinite impulse response (IIR) filter.

16. The speech processing system of claim 13, wherein at least one of the plurality of complex digital filters of the virtual parallel processing chain is characteristic of a gammatone filter.

17. The speech processing system of claim 13, wherein: the plurality of complex digital filters are implemented to perform their digital filtering functions in parallel; and the plurality of complex digital filters are implemented to have overlapping bandwidths that taken together extend substantially over the bandwidth of the digital speech signal.

18. The speech processing system of claim 13, wherein each of the complex digital filters is characterized by a predetermined bandwidth and a predetermined center frequency, the predetermined center frequency of each of the complex digital filters being separated from the predetermined center frequencies of those of the plurality adjacent thereto in the virtual processing chain.

19. The speech processing system of claim 18, wherein the predetermined center frequency spacing between overlapping bandwidths of the complex digital filters is approximately 2%.

20. The speech processing system of claim 18, wherein the predetermined bandwidth of each of the complex digital filters forming the parallel processing chain is 0.75 of its predetermined center frequency.

21. The speech processing system of claim 18 wherein said generating further comprises correcting the estimated instantaneous bandwidth for each one of the digitally filtered signals generated by one of the complex digital filters, said correcting further comprising:

determining a difference between the estimated instantaneous frequency for two of the digitally filtered signals generated by digital filters having bandwidths overlap-

23

ping the bandwidth of the one of the digital filters that generated the digitally filtered signal being corrected; and

dividing the determined difference by the predetermined center frequency spacing.

22. The speech processing system of claim 13, wherein the integration kernel is characteristic of a second order gamma IIR filter.

23. The speech processing system of claim 13, wherein the integrated-product set has at least one zero-lag complex product and at least one two-or-more-lag complex product in place of the at least one single-lag complex product.

24. The speech processing system of claim 13 wherein said generating further comprises integrating the product sets formed for each of the plurality of digitally filtered signals over a predetermined period of time to generate the estimated instantaneous frequency and the instantaneous bandwidth for each of digitally filtered signals.

25. An apparatus for extracting speech content within a digitized speech signal, the speech content being characterized by at least one formant, each of the at least one formants characterized by an instantaneous frequency and an instantaneous bandwidth, the speech signal including a sequence of one or more of the at least one formants, the apparatus comprising:

a reconstruction processor configured by program instructions to receive and operate on samples of the digital speech signal, the reconstruction processor computationally implementing a plurality of complex digital filters, the plurality of complex digital filters implemented to perform their processing in parallel on each sample of the digital speech signal, each of the complex digital filters characterized by a bandwidth that overlaps with the bandwidth of at least one other of the plurality of complex filters, each of the complex digital filters generating as an output one of a plurality of digitally filtered signals, each of the digitally filtered signals comprising discrete values for each sample of the digital speech signal processed, each of the digitally filtered signals including a real component and an imaginary component;

an estimator processor configured by program instructions to receive the plurality of digitally filtered signals from the reconstruction processor, the estimator processor computationally implementing an estimator object, the estimator object being instantiated for each one of the generated digitally filtered signals, each instantiation of the estimator object configured to generate an estimated instantaneous frequency and an estimated instantaneous bandwidth from each of the plurality of digitally filtered signals using a product set formed of each of the plurality of digitally filtered signals; and

a post-processing processor configured by program instructions to receive the estimated instantaneous frequency and instantaneous bandwidth estimates for each of the plurality of digitally filtered signals from the estimator processor, the post-processing processor further configured by program instructions to identify each of the sequence of one or more formants of the digital speech signal as one of the at least one formants based on the received estimated instantaneous frequencies and estimated instantaneous bandwidths of the plurality of filtered signals, the post-processing processor also configured by program instructions to reconstruct the speech content of the digital speech signal using the identified formants.

24

26. The apparatus of claim 25, wherein each instantiation of the estimator object further comprises a computationally implemented integration kernel configured to integrate the product sets formed for each of the plurality of filtered signals over a predetermined period of time to generate the estimated instantaneous frequency and the instantaneous bandwidth for each of filtered signals.

27. The apparatus of claim 26, wherein the integration kernel is characteristic of a second order gamma IIR filter.

28. The apparatus of claim 26, wherein the estimated instantaneous frequency and the estimated instantaneous bandwidth from each of the plurality of digitally filtered signals is generated using a product set formed by the estimator object from each of the plurality of filtered signals in combination with at least one single lag-delay of each of the plurality of digitally filtered signals.

29. The apparatus of claim 26, wherein the estimated instantaneous frequency and the estimated instantaneous bandwidth from each of the plurality of digitally filtered signals is generated using a product set formed by the estimator object from each of the plurality of filtered signals in combination with a two-or-more-lag delay of each of the plurality of digitally filtered signals.

30. The apparatus of claim 25, wherein at least one of the complex digital filters computationally implemented by the reconstruction processor is characteristic of a gammatone filter.

31. The apparatus of claim 30, wherein the predetermined center frequency spacing is approximately 2%.

32. The apparatus of claim 31, wherein the predetermined bandwidth of each of the complex digital filters is approximately 0.75 of its predetermined center frequency.

33. The apparatus of claim 25, wherein each of the complex digital filters includes a predetermined bandwidth and a predetermined center frequency, the predetermined center frequency of each of the complex digital filters being separated from the predetermined center frequencies of those complex digital filters having a bandwidth that overlaps therewith by a predetermined center frequency spacing.

34. The apparatus of claim 33 wherein the estimator processor is further configured to implement a correction process that receives the estimated instantaneous frequency and the estimated instantaneous bandwidth from the estimator processor, the correction process providing a corrected estimated instantaneous bandwidth for each of the filtered signals to the post-processing module using a difference between the estimated instantaneous frequency for two adjacent complex filters in the chain divided by the predetermined center frequency spacing.

35. The apparatus of claim 34 wherein the correction process further provides a corrected estimated instantaneous frequency for each of the filtered signals to the post-processing processor by applying the corrected bandwidth for each of the filtered signals in a best-fit equation.

36. The apparatus of claim 25 wherein the reconstruction processor, the estimator processor and the post-processing processor are implemented as one or more digital processors.

37. The apparatus of claim 25 wherein at least one of the one or more digital processors is a general purpose microprocessor.

38. The apparatus of claim 25 wherein the reconstruction processor, the estimator processor and the post-processing processor are implemented as one or more DSP components.